

IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

FILED
U.S. DISTRICT COURT
EASTERN DISTRICT OF TEXAS

JAN 10 2003

DAVID J. MALAND, CLERK

BY
DEPUTY

1. RAYTHEON COMPANY

Plaintiff,

v.

1. BROCADE COMMUNICATIONS
SYSTEMS, INC.
2. ORACLE CORPORATION
3. OVERLAND STORAGE, INC.
4. QUALSTAR CORP.
5. QLOGIC CORPORATION
6. RICOH CORPORATION
7. SPECTRA LOGIC CORPORATION
8. VERITAS SOFTWARE CORPORATION

Defendants.

CIVIL ACTION NO.

Jury Trial Demanded

2:03cv13-TJW
Ward

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff, Raytheon Company (hereinafter "Raytheon") files this complaint against the following defendants: 1. BROCADE COMMUNICATIONS SYSTEMS, INC. ("Brocade"); 2. ORACLE CORPORATION ("Oracle"); 3. OVERLAND STORAGE, INC. ("Overland"); 4. QUALSTAR CORP. ("Qualstar"); 5. QLOGIC CORPORATION ("QLogic"); 6. RICOH CORPORATION ("Ricoh"); 7. SPECTRA LOGIC CORPORATION ("Spectra Logic"); and 8. VERITAS SOFTWARE CORPORATION ("Veritas"; collectively, "Defendants").

NATURE OF ACTION AND JURISDICTION

1. This is an action for patent infringement arising under the Patent Laws of the United States, 35 U.S.C. § 101 et. seq.
2. Jurisdiction of this Court is proper pursuant to 28 U.S.C. §§ 1331, 1332 and 1338(a).

PARTIES

3. Raytheon is a Delaware Corporation having principal places of business at 141 Spring St., Lexington, MA 02421 and 6620 Chase Oaks Boulevard, Plano, TX 75023.
4. On information and belief, Oracle is a Delaware corporation having a principal place of business at 500 Oracle Parkway, Redwood City, CA 94065. Oracle's Registered Agent for service of process in Texas is Corporation Service Company d/b/a CSC, who may be served at 800 Brazos Street, Suite 750, Austin, TX 78701.
5. On information and belief, Brocade is a Delaware corporation having a principal place of business at 1745 Technology Drive, San Jose, CA 95110. Brocade's Registered Agent for service of process in Texas is CT Corporation System, who may be served at 350 N. St. Paul St., Dallas TX 75201.
6. On information and belief, Overland is a California corporation having a principal place of business at 8975 Balboa Ave., San Diego CA 92123. Overland's Registered Agent for service of process in Texas is Corporation Service Company d/b/a CSC, who may be served at 800 Brazos Street, Austin, TX 78701.
7. On information and belief, Qualstar is a California corporation having a principal place of business at 9621 Irondale Ave., Chatsworth, CA 91311. Qualstar's agent for service of process in Texas is the Secretary of State, who may be served at 1019 Brazos Street, Austin, TX 78071.
8. On information and belief, QLogic is a Delaware corporation having a principal place of business at 26600 Laguna Hills Drive, Aliso Viejo, CA 92656. Qlogic's Registered Agent for service of process in Texas is John W. Warren, who may be served at 12115 Heatherwick Dr., Cypress, TX 77429.

9. On information and belief, Ricoh is a corporation having a principal place of business at 5 Dedrick Place, Caldwell, NJ 07006. Ricoh's Registered Agent for service of process in Texas is CT Corporation System, who may be served at 350 N. St. Paul St., Dallas TX 75201.
10. On information and belief, Spectra Logic, having a principal place of business at 1700 N. 55th St., Boulder, CO 80301, is also or was formerly known as Western Automation Laboratories, Inc., which is a Colorado corporation. Spectra Logic's Registered Agent for service of process in Texas is CT Corporation System, who may be served at 350 N. St. Paul St., Dallas TX 75201.
11. On information and belief, Veritas is a Delaware corporation having a principal place of business at 350 Ellis Street, Mountain View, CA 94043. Veritas' Registered Agent for service of process in Texas is John Beck, who may be served at 15851 N. Dallas Parkway, Suite 500, Dallas, TX 75248.

U.S. PATENT NO. 5,412,791

12. On May 2, 1995, U.S. Patent No. 5,412,791 (the "'791 Patent'"), entitled "Mass Data Storage Library" was duly and validly issued to E-Systems, Inc. Raytheon has subsequently acquired the '791 patent.
13. A copy of the '791 Patent is attached hereto as Exhibit A and is made a part of this Complaint.

DEFENDANTS AND THEIR ACTIVITIES

14. Upon information and belief, Defendants have been and still are infringing, and/or actively inducing the infringement of, one or more of Claims 5-8, 12-15, and 23-28 of the '791 Patent by importing, making, using, selling, and/or offering to sell systems embodying the claimed inventions within the United States, and by selling or offering to

sell such materials within this judicial district and elsewhere in the United States of America, and will continue to do so unless enjoined by this Court.

15. Upon information and belief, Defendants have been and still are contributorily infringing at least one or more of Claims 5-8, 12-15, and 23-28 of the '791 Patent by importing, making, using, selling, and/or offering to sell within the United States a component of a patented combination, consisting of a material part of one or more of the claimed inventions, knowing the same to be especially made or adapted for use in the infringement of the '791 Patent and not a staple article or commodity of commerce suitable for substantial non-infringing use. Defendants' infringing conduct is intentional and unlawful and, upon information and belief, will continue unless enjoined by this Court.
16. Defendants' infringing conduct is unlawful and willful, and upon information and belief, will continue unless enjoined by this Court.
17. Defendants' infringement of the '791 Patent has caused and will continue to cause irreparable injury to Plaintiff for which there is no adequate remedy at law.
18. Plaintiff has suffered damages as a result of Defendants' infringement of the '791 Patent and, will continue to suffer damages as a result of Defendants' continued infringement.

WHEREFORE, Plaintiff prays for a judgment:

- (a) that Defendants have infringed the '791 Patent;
- (b) that permanent injunction be issued against further infringement of the '791 Patent by Defendants, its officers, agents, servants, employees, attorneys and all those persons in active concern or participation with Defendants;
- (c) that Defendants be ordered to compensate plaintiff for the infringement, and to pay enhanced damages, up to treble damages as provided by 35 U.S.C. § 284 for

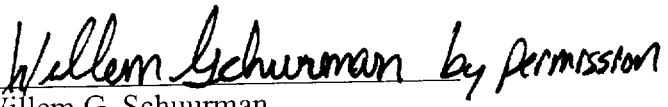
the deliberate and willful nature of its infringement, and to pay interest thereon in the form of both pre-judgment and post-judgment interest;


- (d) that this case is exceptional within the meaning of 35 U.S.C. § 285;
- (e) that Defendants be ordered to pay Plaintiff's costs and expenses and its reasonable attorney fees under 35 U.S.C. § 285; and
- (f) that Plaintiff be granted such other relief as the court deems just and equitable.

JURY TRIAL DEMANDED

Pursuant to Federal Rule of Civil Procedure 38, Plaintiffs demand a jury trial on all issues triable of right by a jury.

Respectfully submitted,


Willem G. Schuurman
State Bar No. 17855200
Attorney-In-Charge



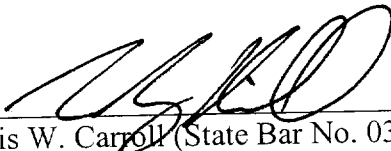
Roderick B. Williams
State Bar No. 21573700
Adam Floyd
State Bar No. 00790699
Christopher V. Ryan
State Bar No. 24037412
VINSON & ELKINS, L.L.P.
The Terrace 7
2801 Via Fortuna Suite 100
Austin, TX 78746
(512) 542-8400 (telephone)
(512) 542-8612 (telecopy)
bschuurman@velaw.com (email)

and

William Dawson
VINSON & ELKINS, L.L.P.
3700 Trammell Crow Center
2001 Ross Avenue
Dallas, TX 75201-2975
(214) 220-7700 (telephone)
(214) 220-7716 (telecopy)

ATTORNEYS FOR PLAINTIFF

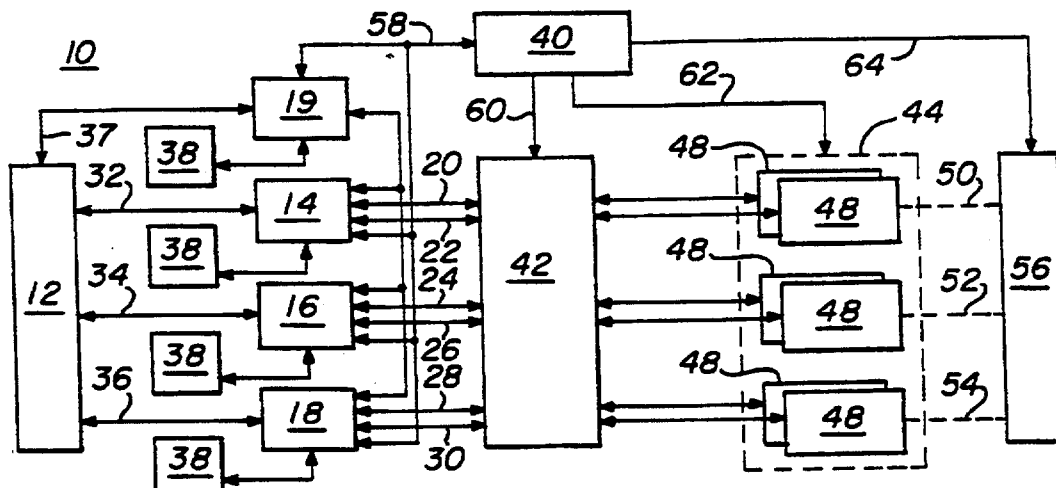
Of Counsel:



Otis W. Carroll (State Bar No. 03895700)
J. Wesley Hill (State Bar No. 24032294)
Ireland, Carroll & Kelley, P.C.
1601 South Broadway, Suite 500
Tyler, TX 75703
(903) 561-1600 (telephone)
(903) 581-1071 (telecopy)

Glenn Lenzen
Raytheon Company.
141 Spring Street
Lexington, MA 02421

[45] **Date of Patent:** **May 2, 1995**

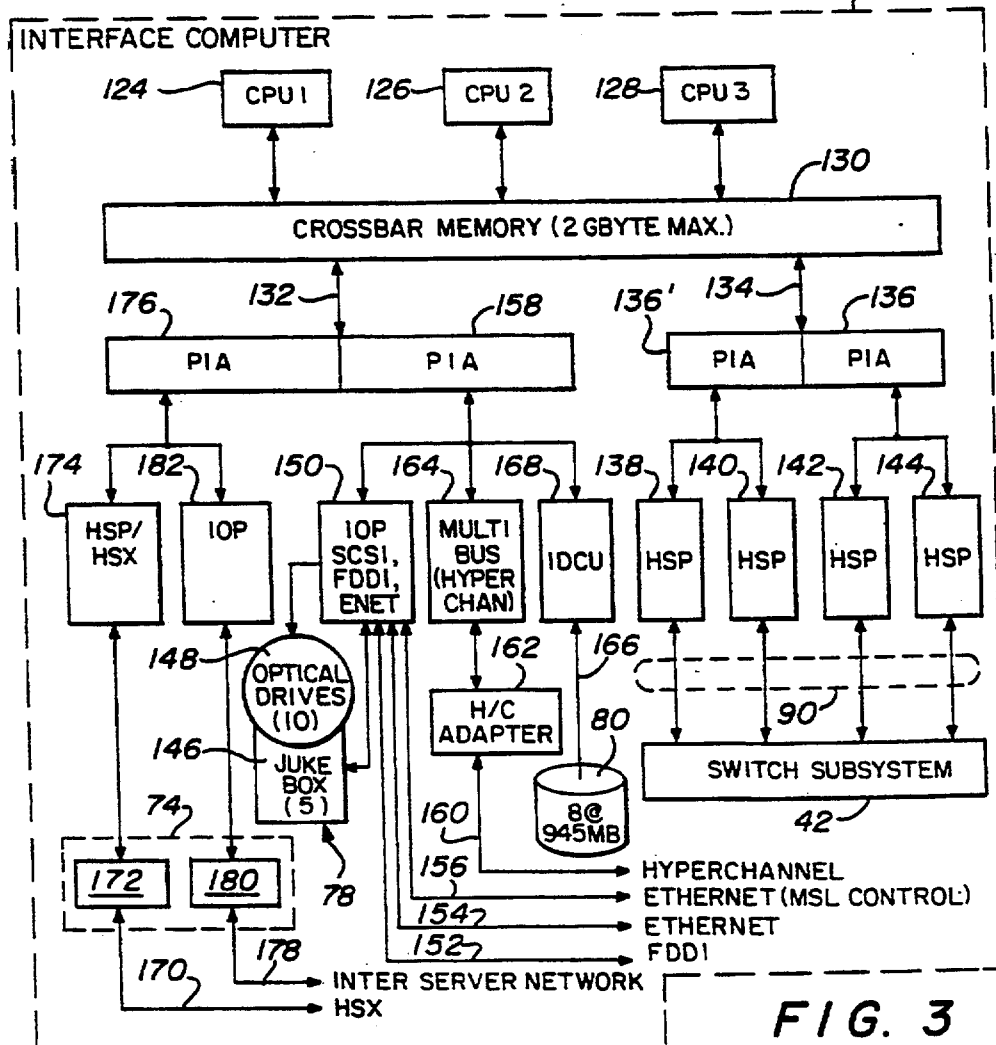



U.S. Patent

May 2, 1995

Sheet 1 of 7

5,412,791

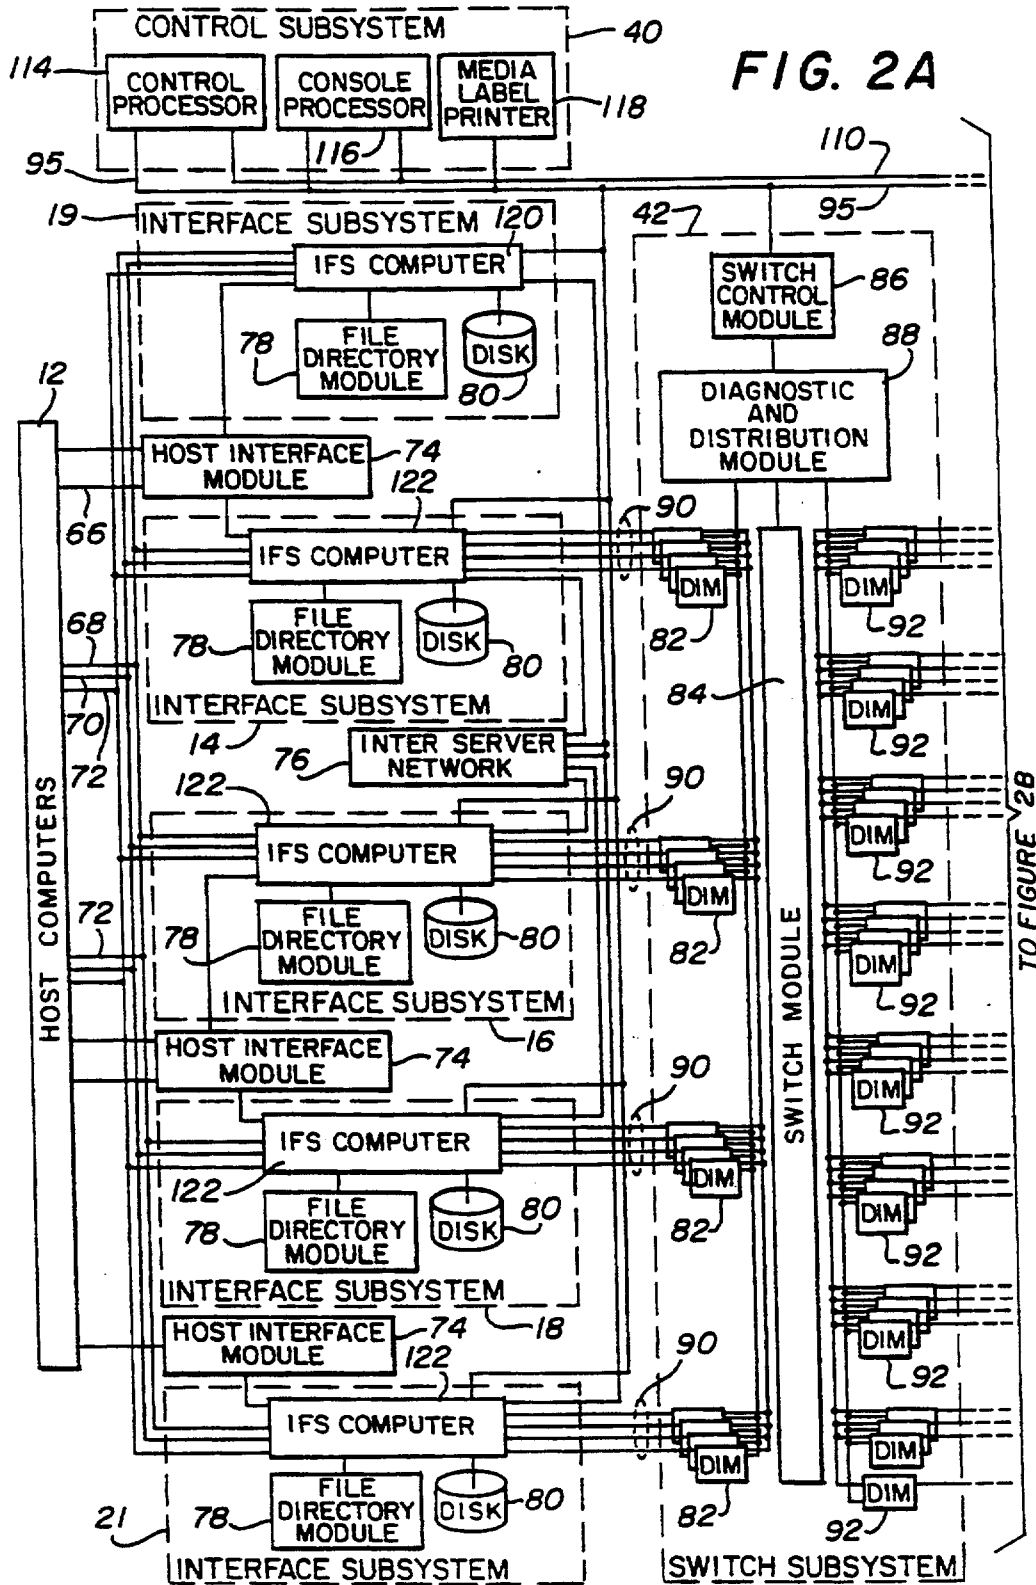


U.S. Patent

May 2, 1995

Sheet 2 of 7

5,412,791




U.S. Patent

May 2, 1995

Sheet 3 of 7

5,412,791

FIG. 2B



U.S. Patent

May 2, 1995

Sheet 4 of 7

5,412,791


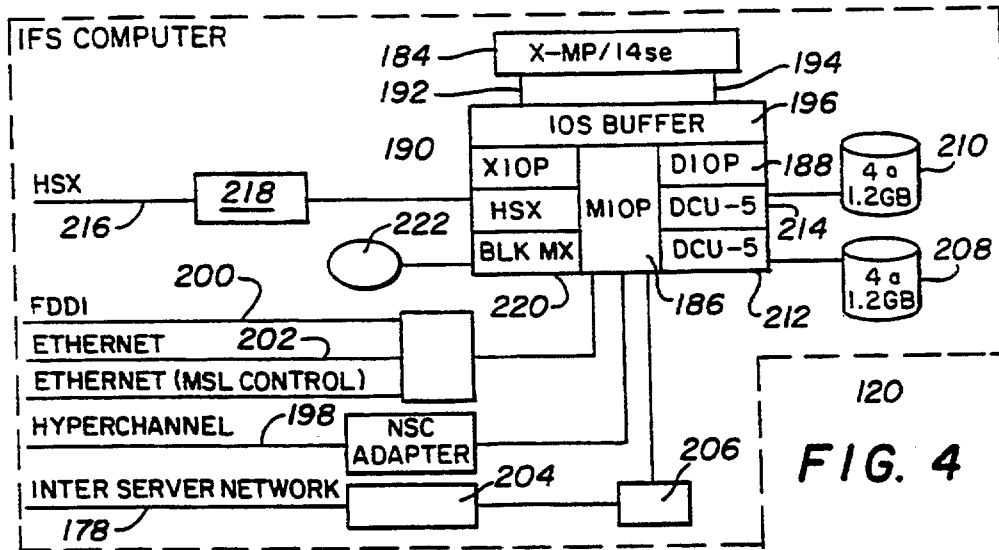


FIG. 5A




FIG. 6




FIG. 9




FIG. 10




FIG. 5B

U.S. Patent

May 2, 1995

Sheet 5 of 7

5,412,791

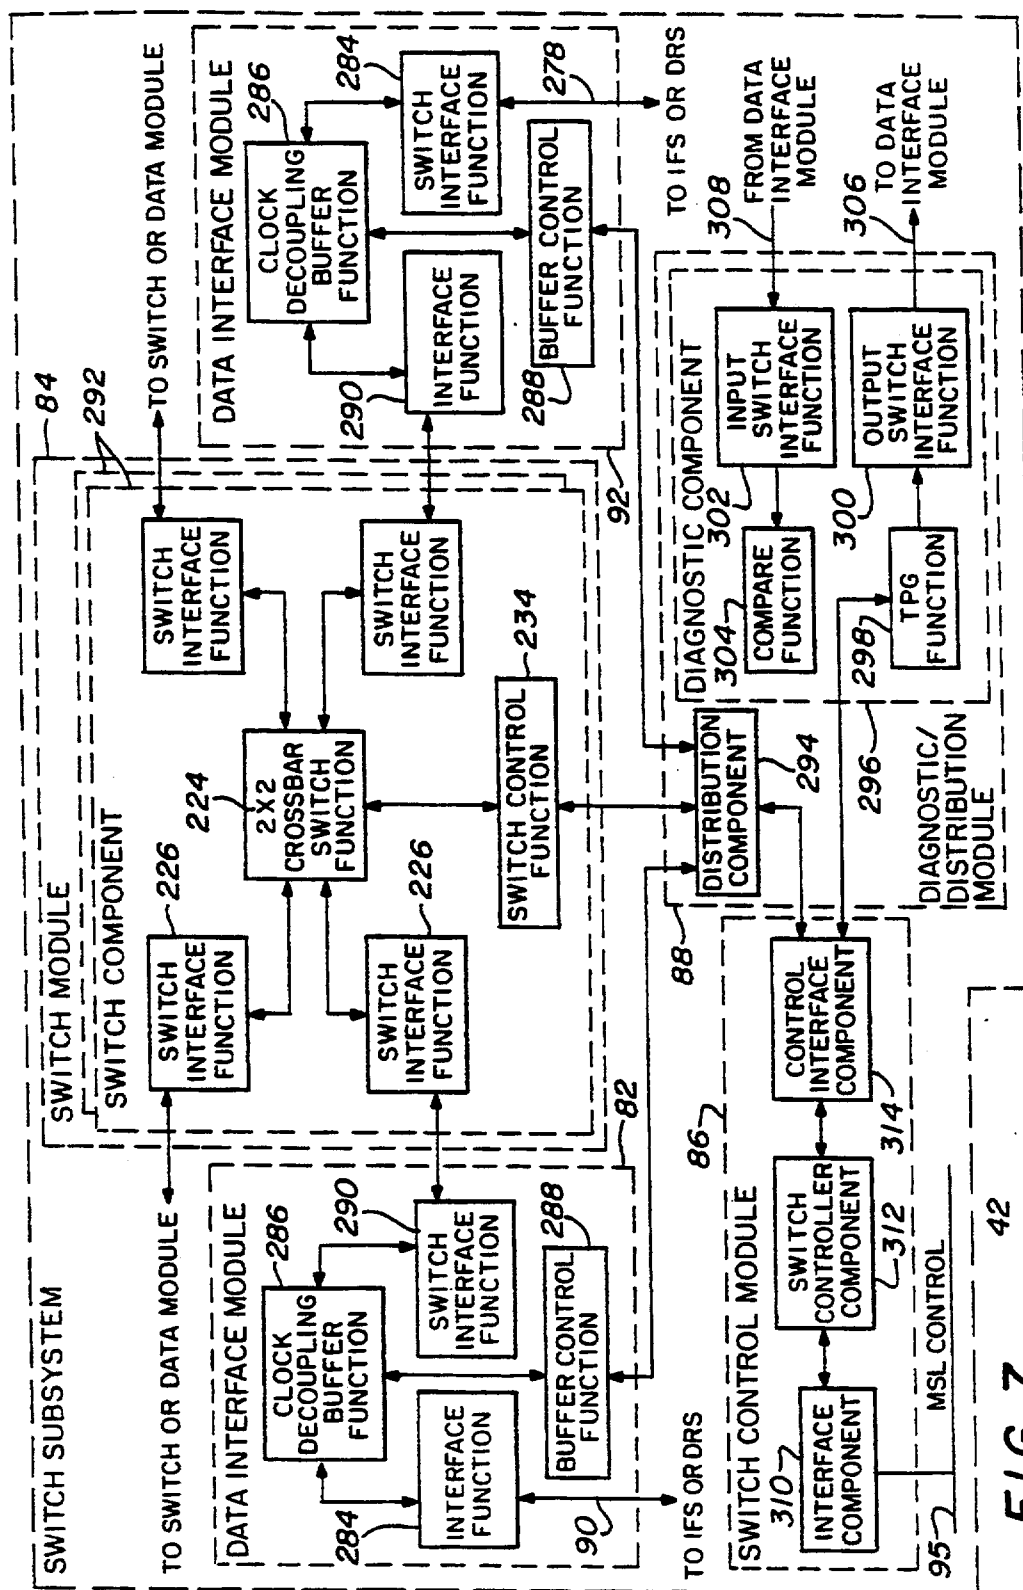


FIG. 7

U.S. Patent

May 2, 1995

Sheet 6 of 7

5,412,791




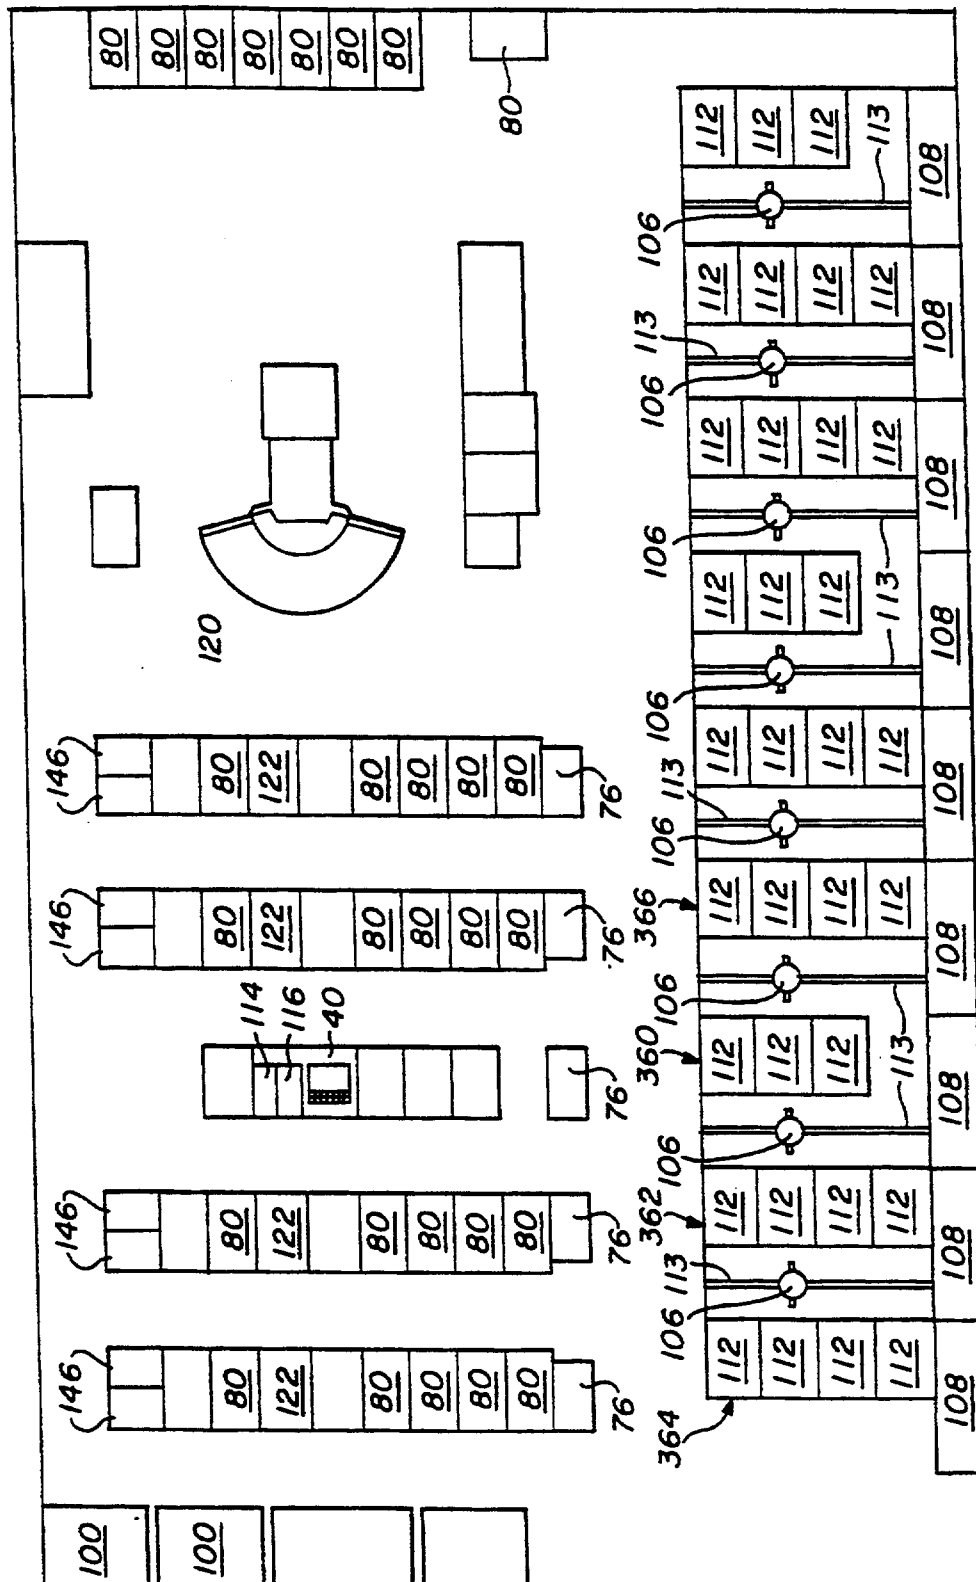
FIG. 8

U.S. Patent

May 2, 1995

Sheet 7 of 7

5,412,791



5,412,791

1

MASS DATA STORAGE LIBRARY

CROSS-REFERENCE TO RELATED APPLICATIONS

This Application is a continuation of prior U.S. application Ser. No. 07/430,134 filed Nov. 1, 1989, now U.S. Pat. No. 5,214,768.

TECHNICAL FIELD

The present invention relates to a mass storage library (MSL) formed in a modular system with control for fast access to store and retrieve digital data and specifically relates to a mass data storage library with modular construction which allows different storage capacities, throughput and access times by optimizing the configuration of standard storage modules, accessors, crossbar switches and interface computers.

BACKGROUND OF THE INVENTION

The present invention relates to a mass data storage and retrieval system formed of one or more arrays of interconnected hardware modules. The need to retain and retrieve large volumes of digital data has resulted in the requirement for ever increasing amounts of data storage devices. Data can be stored, of course, on cassettes, floppy disks, diskettes, hard disks, optical disks, capacitive disks and the like. However, the greater the amount of memory existing, the more difficult it becomes to access the information from multiple input computers in rapid access times and with maximum system fault tolerance.

In order to improve fault tolerance, the prior art has provided disk files comprising a number of stacks having alternate data paths provided using crossbars. This provides some fault tolerance since if one data path fails another can be used.

Other prior art systems utilize a plurality of data storage mediums having a robotic arm control which retrieves the storage disks or cassettes from memory and places them in the recorder. Still other systems include data storage media which are transported to a transducer mechanism such as is utilized for magnetic tape, floppy disk, hard disk and other mass storage devices. After the transducer has completed its use of a particular storage medium, the medium is returned to the storage area from which it came and then another unit is retrieved, transported to and placed in the transducer. Some prior art systems disclose fault tolerant computer systems including a robot maintenance means which is used to replace faulty modules and still others disclose an optical storage/retrieval including a random access arm supplying several users.

Some prior art data storage systems include optical storage wherein a plurality of cartridge storage slots are positioned in diverse locations and orientations and use both vertical and horizontal retrieval systems. Still other systems utilize disk storage and have a disk transport assembly which moves in parallel with the array of the disks independently of the disk retainer and is moveable between a first position operative to have a disk transferred between the array of disks and a disk carrier and a second position operative to have a disk transferred between the disk carrier and the disk retainer.

Still, these systems are inadequate. No matter what drive and storage medium is used in any given installation, use of the system invariably places demands on mass storage that soon exceed the available capacity of

2

the system. Often these demands are not foreseen when the system is initially planned and the need then arises to expand the system. It is difficult to expand such a system in an economical manner. Further, it is difficult in such systems to have parallel operations whereby several computers may be simultaneously accessing the mass storage library where, in some cases, one computer is writing information into the mass storage library while simultaneously another computer is reading information from another part of the mass storage library.

The present invention not only allows the formation of mass storage libraries with a module system to provide versatility, but it also provides the concept of global access. Global access provides maximum system fault tolerance thereby increasing system availability. Such global access in the present invention is provided at many levels, from the interface to the storage medium. Each tape server interface computer is connected to a network to interface with one or more host computers. This provides the ability for each host computer to choose any tape server interface computer. Each tape server interface computer is, in turn, connected to the tape drives through a crossbar switch. The crossbar switch, as used in the present invention, allows global access of multiple computer peripherals to a set of computers. This provides better utilization of the peripherals and provides fault tolerance of both computer failure and peripheral failure. Further, the novel configuration of storage modules, which may include both rotational modules and lateral transfer modules, allows any storage medium in the mass storage library such as a cassette tape, floppy disk and the like to be placed into any one of multiple recorder modules.

The rotational storage modules provide a bank of storage media which, in the preferred embodiment, are in the form of cassettes and form a plurality of rows of aligned racks each of which holds a plurality of the data storage media in a vertical array of parallel slots. Each of the racks has the slots in two opposing sides for holding the storage media such as, for instance, tape cassettes. While it is to be understood that any type of storage medium could be used in the present invention, for purposes of simplicity, such medium will hereafter be referred to as cassettes. An automatic tape cassette accessor is associated with each row of storage racks and with a plurality of data recorder modules for retrieving a designated one of the tape cassettes from a slot in one of the storage racks and inserting it in one of the recorder modules for read and write operations. So that the tape cassette on either opposing side of the rack may be retrieved by the same automatic accessor, each rack may be selectively rotated 180°.

Further, each of the accessors is located between two rows of the rotatable racks such that the tape cassettes on opposing sides of the racks in either row may be retrieved by a common accessor.

Thus, if an accessor on either row fails, the tape cassettes normally facing that accessor can be made available to the functioning accessor on the adjacent row by rotating the storage module. This ability to rotate the tape cassettes can also be used to migrate the tape cassettes to idle recorder modules which results in a load leveling of tasks and a more efficient use of recorder resources.

In addition, a lateral transfer trolley is associated with an end of each row of the rotatable racks for moving tape cassettes between the rows. A cassette tape storage

5,412,791

3

unit is associated with and moveable on the trolley from one row to another, the storage unit having slots for storage of the tape cassettes. The slots on the storage unit are accessible by each automatic accessor so as to enable a tape cassette to be transferred from one row to another.

The subsystem for driving the recorder module has an electronic interface and buffering which is controlled with a micro-controller component that customizes the electronics for the desired computer interface. In addition, the drive subsystem contains special read-after-write processing hardware which enables the drive subsystem to provide a system bit error rate substantially better than the recorder alone can provide and makes screening of tapes unnecessary. This circuit does a bit-for-bit compare of the read-after-write data provided by the recorder. If the compare fails or if the internal error correction passes a preset error threshold, then the data block which is in error, or suspect, is rewritten on the tape. This procedure is novel because the block of data is rewritten on a new block of the tape without stopping or repositioning the tape during the operation. This allows areas of tapes which contain flaws beyond acceptable error correction capabilities to be effectively edited from the tape while data is being recorded. In addition, the magnetic tape cassettes used in the mass storage library are bulk erasable.

To maintain a directory of the data that is contained on the library tapes, the mass storage library concept separates the directory data from the library data by placing the directory data on a high density, random access removable, erasable media (e.g., optical disk or magnetic disk). This allows the directory information to be quickly accessed without having to access the actual library data. The directory media are configured using a "jukebox" system to allow an expandable, on-line archive of the directory data.

The control of the tape drives is divided between network and direct interface depending upon the operational state of the tape drive. When the tape drive controller is in an off-line state, control may be provided by any computer over the network. This allows performance of diagnostics and, in the case of tape drives, loading and unloading, and threading and positioning, to be accomplished by any computer on the network. When the tape drive controller is in an on-line mode, control is provided through the interface in the conventional manner by the on-line interface computer.

SUMMARY OF THE INVENTION

Thus, the present invention relates to a mass data storage and retrieval system comprising a plurality of first data storage modules for containing data storage elements forming a mass information storage library, a data directory archive comprising at least a second data storage medium for maintaining a directory of the library information contained on each data storage module storage element, and a plurality of data recorder modules, each of which is capable of receiving a selected storage element from a data storage module in the mass storage library. A plurality of parallel operating interface computers are coupled to a plurality of host computers for receiving data and for generating request signals to access desired information on the first data storage module storage elements. A directory function is included in the interface computers and the data directory archive for receiving the request signal, locating the address of the storage element containing

4

desired information in the data archive and generating a data location output signal for the storage element. A control computer is coupled to the directory computer and the mass storage library for receiving the data location output signal from the directory computer and in response thereto generating a first signal for causing the storage element in the mass storage library containing the requested information to be loaded into a selected data recorder. A data channel couples the interface computer generating the access request signal to the selected data recorder module for selectively reading data from and writing data to the selected storage elements.

The data channel which couples the interface computers to the data recorder modules comprises a plurality of crossbar switches, each having a plurality of independent input/output ports, any two of which may be selectively connected. One port of each crossbar switch is coupled to one port of an adjacent crossbar switch so as to serially connect the crossbar switches. A switch control module coupled to each crossbar switch generates output routing signals coupling any two selected ports in each crossbar switch to each other. A control computer generates an output signal that is coupled to the switch control module for determining the routing signals coupling two predetermined ones of the ports of each crossbar switch together to couple any interface computer generating an access request signal to any selected data recorder.

The mass storage library comprises a plurality of rows of aligned racks each of which racks holds a plurality of data storage elements in parallel, slots, each of the racks having the slots in two opposing sides for holding the storage elements, an automatic storage element accessor located between and associated with the row of storage racks on each side and with one of a plurality of data recorder modules for retrieving a designated one of the storage elements from a slot in one of the associated racks and inserting it in the associated recorder module for read and write operations, and means coupled to each of the racks so that the storage element on either opposing side of the rack may be retrieved by the associated automatic accessor (i.e., by selectively rotating each rack 180°).

Each of the automatic tape accessors is located between two rows of the rotatable storage racks such that storage elements on opposing sides of the racks in either row may be retrieved by the accessor. A lateral transfer trolley is associated with an end of each row of rotatable racks for moving data storage elements between rows and has a storage unit associated with and movable by the trolley from one row to another, the storage unit having slots for storage of the storage elements. The storage unit is movable by the trolley such that the slots are accessible by each accessor so as to enable a cassette tape to be transferred from one row to another row.

BRIEF DESCRIPTION OF THE DRAWINGS

A more complete understanding of the invention may be had by reference to the following detailed description when taken in conjunction with the accompanying drawings, wherein:

FIG. 1 is a diagrammatic representation of the mass storage library of the present invention;

FIG. 2A and 2B are a block diagram of the mass storage library;

5,412,791

5

FIG. 3 is a diagrammatic representation of the tape server computer;

FIG. 4 is a diagrammatic representation of the disk server computer;

FIG. 5A is a schematic representation of a single crossbar switch node in a first configuration;

FIG. 5B is a schematic representation of a single crossbar switch node in a second configuration;

FIG. 6 is a schematic representation of a plurality of the crossbar switch nodes which may be serially coupled to each other;

FIG. 7 is a diagrammatic representation of the switch subsystem utilizing the crossbar switches;

FIG. 8 is a schematic representation of the drive subsystem coupled to the switch subsystem and which drives the recorder modules;

FIG. 9 is a diagrammatic representation of one configuration of the rotational storage module;

FIG. 10 is a schematic representation of a second configuration of the rotational storage module; and

FIG. 11 is a diagrammatic representation of the floor plan of the present invention in its preferred installation form.

DETAILED DESCRIPTION OF THE DRAWINGS

In the mass storage system 10 shown in FIG. 1, computers 14, 16, 18 and 19 provide the control for communicating with the client users or host computers 12, accepting data from one or more of a plurality of host computers 12 on lines 32, 34 and 36 storing the data and maintaining a media directory or archive 38. Computer 19 utilizes an on-line archive capability while computers 14, 16 and 18 provide direct access to the mass library data storage units.

To facilitate the two types of functions, two types of interface computers are used. The first, computer 19, is of comparable power as a CRAY computer that provides a high performance on-line archive capability. Hereafter, this file server is referred to as an interface subsystem (IFS) disk server. The other interface subsystem, computers 14, 16 and 18, are of comparable power as CONVEX computers that provide a direct tape storage capability and are therefore referred to hereafter as IFS tape servers. The computers 14, 16, 18 and 19 may be accessed simultaneously and in parallel by one or more of the host computers 12. In like manner, multiple outputs are provided by the IFS tape server computers 14, 16 and 18 as illustrated by lines 20 and 22, 24 and 26 and 28 and 30. It is to be understood that two output lines are shown for purposes of illustration only and, as will be seen hereafter, more than two outputs from each of the computers 14, 16 and 18 may be operating in parallel and simultaneously.

The outputs of the IFS tape server computers 14, 16 and 18 on lines 20-30 are coupled to a switch module 42 which provides any IFS tape server computer 14, 16 or 18 the ability to be connected to any drive subsystem 48 in drive unit 44. The data storage modules, for example cassette tapes, for the drive subsystem 48 are contained in a transport subsystem 56 and are automatically accessed and loaded into the appropriate drive subsystem 48 by robotic means as illustrated by dashed lines 50, 52 and 54.

Although the description of the invention will be made referring to tape cassettes as the storage elements, it should be understood that the invention contemplates utilization of other and additional storage media.

6

The control subsystem 40 provides control for the allocation and de-allocation of common resources for the mass storage library system. When an interface tape server 14, 16 or 18 or the interface disk server computer 19 receives a command to read or write data, it first requests resources from the control subsystem 40. Computer 40 will initialize and position the appropriate resources and informs the requesting IFS when the resources are available. Control of the recorder resources is then passed to the requesting IFS. Once the operation is complete, the controlling IFS informs control computer 40 that the operation is complete and the control computer 40 de-allocates those resources. Thus, the control computer 40 communicates with the IFS tape servers 14, 16 and 18 and the IFS disk server 19 through line 58 which could be any one of a number of commercially available networks. It also communicates with the switch subsystem 42 through line 60, the drive subsystem 44 through line 62 and the storage/transport subsystem 56 through line 64. These lines 58, 60, 62 and 64 are not necessarily independent.

A more detailed diagram of the mass storage library 10 is illustrated in FIG. 2. The same numerals are used in FIG. 2 for the major system components as were used to represent such components in FIG. 1. FIG. 2 illustrates five interface subsystems 14, 16, 18, 19 and 21 instead of the four shown in FIG. 1. The units 14, 16, 18 and 21 are IFS tape servers while the unit 19 is an IFS disk server. The five interface subsystems are shown for example only and [may,] more or less, be used in the invention. The primary function of the five interface subsystems is to send and receive commands and data to and from the host computers 12. In this example, each interface subsystem (IFS) is connected to a plurality of the host computers 12 through four well known types of external communication networks, the HSX or high speed interface line 66, the Ethernet channel on line 68, the FDDI on line 70 and the HYPERchannel® on line 72. Each of these four external networks is coupled to the IFS subsystems through a host interface module (HIM) 74. The host interface module 74 provides the connectivity between the IFS computers 14, 16, 18, 19 and 21 and the network. In the case of the HSX interface on line 66 where it is coupled to the IFS tape servers 14, 16, 18 and 21, an existing Ultra Network (shown in FIG. 3 and FIG. 4) is used to provide the HSX/HSP conversion. The IFS disk server 19 communicates with any IFS tape server 14, 16, 18 and 21 when it needs to store or retrieve data from the cassette tape library via the inter-server network 76.

Each IFS unit 14, 16, 18, 19 and 21 has a file directory module (FDM) 78 which comprises at least one removable storage medium for maintaining a directory of the information contained on each data storage module. In the case where the FDM 78 is used in the IFS disk server 19, it is a proprietary unit that maintains file directory information for the disk server 19. The FDM 78 contains information indicating if the requested file is on disk or has been archived to tape. If the FDM 78 is configured in an IFS tape server 14, 16, 18 or 21, then it includes a hardware configuration that comprises an optical disk jukebox which contains erasable optical disk drives and storage modules. Each of the IFS units 14, 16, 18 and 21 also has a disk drive module 80 which provides that IFS with the capability of storing file directory information and file data so that in the case that certain data needs to be accessed often, that information can be stored in the disk drive module 80 from

5,412,791

7

the cassette in the library, thus avoiding the requirement that the cassette in the tape library be accessed frequently. In the case of the IFS disk server 19, the disk drive 80 also contains both file directory information and file data. When the disk drive 80 reaches a predetermined fill level, the resident software spools the data off to an IFS tape server 14, 16, 18 or 21 via the INTER SERVER NETWORK 76. The IFS 19 also has the ability to request that its data be stored off to cassette tapes in the mass storage library. In that case, the resident software performs the function automatically. IFS 19 is still responsible for providing that data and file information to the user even though the data has been stored on cassette tape. IFS 19 has the capability to store data in any of the IFS tape servers 14, 16, 18 or 21. When the disk drives 80 are used in the IFS tape server configuration, they provide storage for system software such as system administration and diagnostic routines. The disk drives 80 also store data necessary to access further file directory information stored on optical disk.

The switch subsystem 42 comprises four modules: the data interface module 82 (DIM); the switch module 84; the switch control module 86; and the diagnostic/distribution module 88. The switch subsystem 42 provides the IFS tape servers 14, 16, 18 and 21 with the ability to be connected to any drive subsystem 48 in drive unit 44. This connection is performed by providing a one-for-one connection between any of the four outputs or interfaces 90 from each IFS tape server 14, 16, 18 and 21 to any of forty (nominal) drive subsystems 48 in drive unit 44. The data interface module 82 (DIM) provides the one-for-one connection to the interfacing subsystems on both sides of the switch subsystem 42. The DIM 82 in the switch subsystem 42 contains the functions necessary to provide data to switch module 84 that selectively couples that data to another DIM 92 on the other side of the switch module 84. The switch module 84 is composed of modular crossbar switch functions that provide connectivity between all IFS DIMs 82 to all drive subsystem DIMs 92 and vice-versa.

The switch control module 86 in switch subsystem 42 interfaces with the mass storage library (MSL) through the control Local Area Network (LAN) channel 95 (line 60 in FIG. 1). On channel 95, it receives commands and returns subsystem status. The switch control module 86 receives commands from the control subsystem 40 and initiates either a switch function or commands the diagnostic/control module 88 to run a diagnostic test. Prior to any read or write operations, the control subsystem (CNS) 40 commands the switch subsystem 42 to connect any IFS DIM 82 to any drive subsystem DIM 92. All subsystem status is routed to the switch control module 86 which provides the control to format and send status to the CNS 40. The switch module 84 incorporates a crossbar switch component which will be disclosed in more detail hereafter and which connects any IFS interface 90 to any drive subsystem 48 in drive unit 44.

The diagnostic/distribution module 88 provides the central point for distributing control information within the switch subsystem 42. The diagnostic/distribution module (DDM) 88 has a distribution component that distributes control information from the switch control module 86 to all other modules. DDM 88 also provides a diagnostic function that generates and compares test data. The DDM 88 is commanded by the switch control module 86 to generate and send diagnostic data through the subsystem.

8

The drive unit 44 has a multiplicity of drive subsystems 48, each of which comprises a drive controller module 94 (DCM) and a recorder module 96 (RM). The main function of the drive subsystem 48 (DRS) is to record data on and playback recorded digital data from cassettes in the MSL. DRS 48 interfaces with the switch subsystem 42 (SWS), the transport subsystem 56 (TRS) and the control subsystem 40 through the control LAN 95. The DRS 48 interface with the SWS 42 and the control LAN 95 is an electrical interface while the interface to the TRS 56 is a mechanical interface in which storage media are retrieved or inserted by a robotic function. The DCM 94 interfaces with the MSL control LAN 95 to receive commands from and return subsystem status to the CNS 40.

When a read or write operation is to be performed, the commanded IFS 14, 16, 18 or 21 requests resources from the CNS 40. CNS 40 responds with a resource allocation list that contains the allocated drive subsystem 48. The allocated DRS 48 is commanded throughout the read/write operation by that IFS 14, 16, 18 or 21 via a data channel through the SWS 42. Once the operation is complete, status is returned from the DRS 48 to the controlling IFS 14, 16, 18 or 21.

The DCM 94 internally interfaces with the recorder module 96. THE DCM 94 commands the RM 96 to perform various recorder functions such as start/stop, thread tape, record and the like. The RM 96 returns status to the DCM 94 upon completion of a command. The current RM 96 is a high speed data recorder capable of accepting and sending data at the rate of up to 240 Mbps. As technology improves, this configuration allows adaptation to a higher speed RM 96. The RM 96 interface is to the DCM 94 for command/status and data and also interfaces to the TRS 56 via a robotic accessor. A recorder module 96 can be loaded either manually or by a robotic accessor. The tape support subsystem 98 (TSS) provides system level support and repair for the media. All of the TSS modules 98 are stand alone modules that require manual loading. Degasser module 100 provides the system with the ability to bulk erase the tapes. Tape repair module 102 provides the system with the ability to repair broken tape or cassettes. Winder/cleaner module 103 provides the system with the ability to clean the tape media and re-pack the media.

The tape transport subsystem 56 (TRS) moves tape cassettes between the drive subsystem 48 and the storage subsystem 104 (STS). The primary function of the TRS 56 is to retrieve cassettes from the recorder modules 96 and store them in the STS 104 and vice-versa. The transport subsystem 56 (TRS) in addition to the STS 104, also includes the cassette accessor module 106 and the lateral transfer module 108. All of the TRS 56 modules are commanded by the CNS 40 via the MSL storage LAN 110 (line 64 in FIG. 1).

The cassette accessor modules 106 (CAM) are implemented by any well known robotic function. The CAM grasps a cassette tape from a storage module 112 and inserts it either into a RM 96 in a DRS 48 or in a slot in the lateral transfer module 108 (LTM) or in another storage module 112 in STS 104. The cassette accessor module 106 traverses an isle servicing any number of drive subsystems 48 (although only four are shown in each of the eight groups in the drive unit 44 in FIG. 2 for purposes of simplicity of the drawings) and any number of storage modules 112 illustrated in the storage subsystem 104. The total number of storage modules in

5,412,791

9

an isle is dependent on the location of the isle as will be seen hereafter. The cassette accessor module moves in both a horizontal and vertical axis. The accessor is configured with up to four independent arms and rotates about the horizontal axis.

The LTM 108 provides the TRS 56 with the ability to move tapes between storage rows with a minimum amount of user interaction. The LTM 108 spans the width of the transport subsystem 56 and is located at the rear of the isles as can be seen more clearly in FIG. 11. The LTM 108 comprises sectional transfer components that are connected end-to-end and are stacked two units high.

The storage subsystem 104 (STS) provides the storage modules 112 for accommodating cassette tapes. The exact number of cassettes accommodated is dependent on the STS configuration. The storage modules 112 are rotational modules (RTM) and are stacked two high. Each module can hold eighty large cassette tapes or forty medium tapes. (See FIGS. 9 and 10.) The storage modules 112 rotate 180° about the vertical axis in 90° increments and are commanded by the CNS 40 via the MSL storage LAN 110. The rotational modules 112 provide each of the CAMs 106 with the ability to access up to four rows of cassette tapes. This ability improves system availability and reliability because an accessor can access tapes from an adjacent aisle when the storage module is rotated. This is especially useful where all of the recorder modules 96 available to one CAM are busy or where a cassette accessor has become disabled. The LTM 108 supports load leveling by transferring cassettes between any two aisles.

The CNS 40 includes a control processor 114, a console processor 116, a high speed printer (not shown), and a media label printer 118. The control processor 114 and the console processor 116 are Sun 3 Series 200 workstations. They share a disk pool and provide immediate mutual redundancy. Both processors are connected to the line printer (not shown) that is capable of printing at least 300 lines per minute. Both processors are also connected to a media label printer 118 through the MSL control LAN 95. The media label printer 118 produces machine readable and human readable media labels. This configuration enables the CNS 40 to tolerate the failure of a single disk drive with no degradation in performance or throughput. In the event of a processor failure, either processor can be configured to perform the entire processing function.

As stated previously, the interface subsystem (IFS) provides the capability for the mass storage library 10 to send and receive data to and from other computers over a variety of computer networks. The IFS is composed of a set of loosely coupled modules that provide the flexibility required to meet the needs of the user complex. Interchangeable and expandable IFS components are provided so that each IFS module can be configured to support specific performance and functional needs. Each IFS uses commercially available hardware and software products. Each IFS is composed of an IFS computer module (ICM) and a file directory module 78 (FDM) and is coupled to a host interface module 74 (HIM). The IFS disk server 19 includes an IFS computer module 120 (ICM) which is a CRAY X-MP/14se Super Computer to provide a high performance on-line archive capability. Each of the other IFS tape server units 14, 16, 18 and 21 utilize an ICM 122 which is a CONVEX 232 Computer and this implementation provides the user with a direct interface to the mass storage

10

drives. The IFS tape server computers 122 and the IFS disk server computer 120 perform the same basic function but offer different performance features. The data and communication length between the IFS tape servers and the IFS disk server is provided by the inter-server network 76.

The host interface module 74 (HIM) contains the physical interface components that provide the system with the capability to communicate with the low speed and high speed user system interfaces. It is a logical grouping of the components and adapters, well known in the art, that provide the capability to interface with the Ethernet Network on line 68, the HYPERchannel® network on line 72, the FDDI token ring network on line 70 and the high speed HSX channel on line 66. Thus, the HIM 74 enables the system to communicate with user systems that are connected to one or more IEEE Standard 802.3 Local Area Networks (LANs). HYPERchannel® processor adapters provide the system with the capability to communicate with the user systems that are connected to one or more HYPERchannel® networks. Two different HYPERchannel® processor adapters are required for the system. The A130 HYPERchannel® adapter provides the capability for the CRAY Computer 120 to communicate with the user systems that are connected to a HYPERchannel® network. This adapter is specifically designed with a pair of 16-bit high speed asynchronous interfaces for connectivity to the IFS Computer 120. An A400 HYPERchannel® adapter combined with the PL150 processor interface provide the capability for the CONVEX C232 Computer 122 to communicate with the user systems that are connected to a HYPERchannel® network. These items are all commercially available and are well known in the art.

The fiber-distributed data interface (FDDI) utilizes a SUPERNET chip-set that implements the FDDI standard.

A direct HSX channel connection is made to the CRAY IFS Computer 120. The HSX channel is capable of sustaining an effective transfer rate of at least 200 Mbps through the HSX data link.

The CONVEX C232 Computer 122 does not provide an I/O channel that is directly compatible with the CRAY HSX channel. Therefore, Ultra products are utilized to provide connectivity between the CONVEX C232 Computer and the HSX channel. The system is capable of achieving 200 Mbps between the CONVEX C232 Computer 122 and the Ultra Network Hub.

The CONVEX C232 Computer 122 is disclosed in detail in FIG. 3. It contains three CPUs 124, 126, and 128. Each of the computers 124, 126, and 128 is coupled bidirectionally to a crossbar memory 130 which has two I/O channels 132 and 134. The use of a multi-processor design allows for the run time software to allow dissimilar events or multiple interrupts to be processed in parallel. The Computer 122 uses a 64-bit word and is capable of having two GBytes of physical memory attached. This memory is partitioned into a system memory and a tape buffer memory. The output of the crossbar memory 130 on I/O channel 134 is coupled through peripheral interface adapter 136 to four high speed parallel interface units 138, 140, 142 and 144. The output of these high speed parallel interface units are the outputs 90 shown in FIG. 2 from the IFS Computer 122. They are used to transfer data to the Switch Subsystem 42. The tape drive buffer forming part of the crossbar memory 130 is used to match the tape drive data transfer rate

5,412,791

11

with the transfer rate of the connected interface to the user system. The tape buffer also provides for more efficient use of the drives. Data transfers to the drives occur when enough data has accumulated in the buffer to make an efficient transfer to the tape drive. Conversely, data transfers to the user system occur when the buffer has accumulated enough data from the tape drive to make an efficient transfer. The entire two GByte memory range in the crossbar memory 130 is available to each Computer 124, 126, 128 and the I/O channels 132 and 134. All four of the tape drives through the high speed parallel interface unit 138, 140, 142 and 144 can be operating simultaneously. Each CONVEX ICM 122 has up to five file directory modules 78 (FDM) each consisting of a juke box 146 for storing optical disks and two optical drives 148. The FDM 78 provides storage for the file system directories. To provide fast and frequent access to directory information, the directory data is on media separate from the actual data. By querying the FDM 78, users access the directory information of a cassette tape without mounting a tape into a drive. This allows the tape drive to be used more effectively by only handling data read and write operations. Tapes can be removed from the cassette library while directory data is maintained. When information on tape is exported to other sites, the removable media serves as a means to transport directory information to the receiving site, eliminating the need for an exhaustive survey of delivered tape media.

To achieve goals of large storage capacity, removability, and fast and frequent access for storing directory information, a juke box containing erasable optical disks in storage bins and erasable optical disk drives are used. These disks are removable from the juke box for archive purposes and may be erased. Each jukebox contains storage for 24 erasable optical disks and space for two optical disk drives. Each of the disks is double-sided and holds a total of one GByte of data. The disk drives 148 use the magneto-optic method of erasability. These products are all currently commercially available. The five FDMs 78 are connected to an input/output processor 150 (IOP) along with the FDDI user input on line 152, the Ethernet user on line 154, and the MSL Control LAN on line 156. The IOP 150 interfaces with the processor interface adapter 158 which is coupled to the I/O channel 132. The user input on HYPERchannel® 160 is coupled through an adapter 162 to interface 164, the output of which is also in communication with the interface adapter 158. The user HSX channel on line 170 is coupled to Ultra circuit 172 in HIM 74 which is in communication with processor interface 174 and from thence to the interface adapter 176, the output of which is coupled to I/O channel 132.

The inter-server network 76, shown in FIG. 2A, is connected through line 178 to an Ultra network hub 180 in host interface module 74 and IOP 182 which allows for communications between the CRAY Computer 120 and the CONVEX Computers 122. The communications consist of files transferred from the CRAY Computer 120 disk storage to be archived on any of the tape drives which are connected to the CONVEX Computer 122. The Ethernet LAN line 156 is used for sending requests to the control subsystem 40 (shown in FIGS. 1 and 2) and for receiving status on the requests from the control subsystem 40.

The CRAY ICM Computer 120 is shown in detail in FIG. 4. The mainframe 184 is a CRAY X-MP/14se which is a vector processor having an estimated perfor-

12

mance of 32 million instructions per second (MIPS). The I/O subsystem (IOS) consists of three different I/O processors (IOPs). The first is a master IOP (MIOP) 186, the second is a disk IOP (DIOP) 188, and the third is an auxiliary IOP (XIOP) 190. Two 100 MByte channels 192 and 194 connect the mainframe 184 to the IOS through the IOS buffer 196.

The MIOP 186 is responsible for providing interfaces to user HYPERchannel® 198, the FDDI channel 200, the Ethernet channel 202, and the inter-server network channel 178. An Ultra network hub 204, which is a part of the host interface module 74, is coupled to the MIOP 186 through a VME interface 206.

The DIOP 188 is responsible for interfacing the mainframe 184 central processor with the disk drives 208 and 210. The disk drives 208 and 210 provide a total of 9.6 GBytes of disk storage using eight disk drives. The drives 208 and 210 are used to provide the online capacity for files stored by the mainframe 184. The drives 208 and 210 interface to the DIOP 188 through controllers 212 and 214.

The XIOP 190 is responsible for providing the HSX interface to the user line 216 through the Ultra interface 218, which consists of an Ultra network hub and link adapter which can extend the HSX channel to a maximum of two kilometers. This interface is needed when the users with an HSX interface are more than 50 feet from the mainframe 184. There is also an IBM Block Mux channel 220 for attaching an IBM tape drive 222 to the mainframe 184. The IBM tape drive 222 is provided for maintenance purposes and as a way for loading new or modified systems software.

As stated earlier in relation to FIG. 2, the switch module 42 is composed of one or more switch components. FIGS. 5 and 5B are schematic representations of one crossbar switch node. The 2×2 crossbar switch 224 is the heart of the switch subsystem 42. The 2×2 crossbar switch contains four independent ports 226, 228, 230, and 232. The switch 224 allows any two of the ports to be connected, thus providing the capability to support up to two simultaneous data transfers. The port independence allows the 2×2 crossbar switch 224 to be reconfigured without interrupting data transfers through the other ports. Thus, any one of the IFS computers may be connected not only to any one of the drive subsystems 48 but may also be connected to each other by switch 224.

Initially, every cross-bar switch node is set in a "straight-through" configuration, as shown in FIG. 5A. Port 226 is connected to port 230, and port 228 is connected to port 232. For a data transfer to occur, these ports connections are opened, and the desired ports are connected to provide the desired connectivity. In FIG. 5B, port 226 is connected to port 232 to provide connectivity between an interface computer (IFS2 and a drive subsystem 48 (DRS 1) (see FIG. 6). Note that switch node 236 need not be configured since it is still in the initialized "straight-through" configuration and therefore provides connectivity to the drive DRS 1. Thus, for any data transfer to take place, only one switch node needs to be configured.

The Non-Blocking algorithm that controls the crossbar switch configuration is based on the premise that each IFS has exclusive access to the vertical ports of the switch nodes in its column. In FIG. 6, IFS 2 has exclusive use of vertical ports 226 and 230, and vertical ports on nodes 254 through "SNY" whereas ports 232 and 228 can be used by a second interface computer (IFS 1)

5,412,791

13

to provide connectivity to DRS 1. This scheme of only allowing IFSs vertical connectivity through its own switch node columns prevents an IFS from interfering or "blocking" any other IFS's data transfer.

Thus, as can be seen in FIG. 6, as an example only, there are three columns each represented as being accessed by an interface computer (IFS 1, IFS 2 and IFS 3) with N switches in each column. The example also includes N rows, each row having an output to a drive subsystem (DRS 1 through DRS N). An input to port 226 of the first 2x2 switch 224 may be coupled to any one of the other ports such as port 230. That port may be connected to any one of the ports of a second 2x2 crossbar switch 254 such as, for example, port 238. I/O port 238 may then be connected to any of the other ports of switch 254 such as port 250. In like manner, the connection may extend from port 250 of switch 254 to port 252 of switch SNY, and through port 256 to port 258 of switch SNZ. From port 260 of switch SNZ, the signal may be coupled to a drive system (DRS N). Any number of desired switches may exist between switch 224 and switch SNZ. Thus, it can be seen that with connections as explained between the switches shown in FIG. 6, any IFS tape server computer 122, represented as IFS 1 through IFS 3 in FIG. 6, may have any one of its outputs 90 (see FIG. 2) coupled to any one of the input/output ports of the 2x2 crossbar switches with the resulting connection being made to any of the drive subsystem units 48 and the drive unit 44. It is in this manner that any IFS computer output could be connected to any particular recorder module 96 to read data from or write data to any cassette tape.

As stated previously, the purpose of the switch subsystem 42 is to provide a logical and physical interface between interface subsystems 14, 16, 18, 21, and 19 (see FIG. 2) and the drive subsystems 48 and the drive unit 44, and to enable an IFS to have global access to all DRS's 48. The SWS 42 provides dynamic connectivity so that IFS-to-DRS reconfigurations occur without interrupting other channels. The SWS 42 is composed of modular components to accommodate all possible mass storage library configurations. This flexibility enables the mass storage library to be configured with a variable number of drives. The independence of each node port enables configured drives to transfer data concurrently and independently to/from separate users. The SWS 42 enables each IFS 14, 16, 18 or 21 to have access to any DRS 48 in the drive unit 44. This global DRS access capability allows the mass storage system to spread the user job load among drives and to minimize the need for media transport since nearer drives are considered prior to distant drives.

As illustrated in FIG. 7, the SWS 42 is composed of four modules. The data interface module (DIM) 82 provides the physical interface to either an IFS 14, 16, 18, or 21 or to a DRS 48. However, for purposes of simplicity, FIG. 7 will be discussed as if DIM 82 were connected to any IFS through line 90 while DIM 92 will be discussed as if it were connected on line 278 to a DRS 48. Again, however, it should be kept in mind that the SWS-42 could function in the reverse manner equally well. Thus, in FIG. 7, the DIM 82 will be considered as providing the physical interface to the IFS, although it could also be to a DRS.

The switch module 84 (SM) is the switching mechanism that contains the 2x2 crossbar switch 224 shown in FIGS. 5A, 5B and 6. The switch control module 86 (SCM) provides an interface to the MSL Control LAN

14

95 which is connected to the CNS 40 (see FIG. 2). SCM 86 receives commands from the control subsystem 40 (CNS) and controls all modules and components of SWS 42. The diagnostic/distribution module 88 (DDM) contains the built-in test equipment for the SWS 42 and distributes control to each SWS 42 component.

As stated earlier, the DIM 82 is the component that provides the SWS 42 external interface to any IFS while DIM 92 provides the external interface to the DRS 48. The interface function 284 in both DIM 82 and DIM 92 contains the balanced differential drivers and receivers specified for the high speed interface. The clock de-coupling buffer unit 286 in both DIMs 82 and 92 provides a means to synchronize the input data with the SWS master clock signal (not shown). The de-coupling is accomplished using memory buffer unit 286 that operate in a ping/pong manner which supports simultaneous memory reads and writes. While input data is written into one memory buffer with the input clock signal, data is read out of the other memory buffer with the SWS 42 master clock signal. The buffer control unit 288 in each DIM controls the ping/pong memory architecture, configures the DIM 82 (or 92) for diagnostic loop-back tests, and provides a control and status interface to the diagnostic/distribution module 88 (DDM). The switch interface unit 290 in each DIM 82 and 92 provides a common interface to the switch module 84. Each DIM 82 and 92 is essentially an extension of the IFS Data Server High Speed I/O channel.

The switch module 84 (SM) is composed of one or more switch components 292. Each switch component 292 contains a 2x2 crossbar switch 224, a switch interface unit 226 coupled to each of the four independent I/O ports and a switch control unit 234. The switch interface unit 226 provides the SWS 42 with modularity by providing connectivity to either a DIM 82 or 92 or to another switch component 292. Multiple switch components 292 may be connected together to achieve the desired number of external SWS 42 interfaces.

As stated earlier, in relation to FIGS. 5A, and B, the 2x2 crossbar switch 224 is the heart of the SWS 42. The 2x2 crossbar switch 224 contains four independent I/O ports. This crossbar switch 224 allows any two ports to be connected, thus providing the capability to support up to two simultaneous data transfers. The port independence allows the 2x2 crossbar switch 224 to be reconfigured without interrupting data transfers through the other ports. The switch control unit 234 provides an interface to the distribution component 294 in the DDM 88. Control and status information are transmitted across this interface. Reconfiguration commands are received by the switch control unit 234 and distributed to the individual crossbar switch ports. Upon command, the switch control unit collects status from the switch component 292 and transmits the status to the distribution component 294. The switch component 292 provides an aggregate I/O bandwidth of 1280 Mbps.

The DDM 88 contains the distribution unit or component 294 and a diagnostic unit 296. The distribution unit 294 provides a command/status interface and distribution for the switch control module 86 (SCM). Upon receipt of a command, distribution unit 294 distributes the control information to the appropriate SWS 42 component. The diagnostic unit 296 includes the SWS 42 built-in test equipment. It contains a test pattern generator 298, an output switch interface 300, an input

5,412,791

15

switch interface 302, and a compare unit 304. Test pattern generator 298 generates known diagnostic test patterns upon command from the SCM 86 and are output through switch interface unit 300 on line 306. Line 306 connects the diagnostic unit 296 to a switch interface function port 226 in switch component 292 through a DIM 82. The diagnostic test patterns are received from the units being tested on line 308 to input switch interface 302 for verification. The compare function 304 verifies the diagnostic test pattern. The diagnostic unit 296 provides the SWS 42 with the capability to generate a diagnostic test pattern, route the test pattern through each, or selected, SWS 42 components, and verify the integrity of the test pattern. Such error detection capabilities of each of the SWS 42 components allows faults to be identified at the lowest replaceable unit level. It provides swift diagnosis of faulty components, and the novel connectivity and dynamic switching capabilities of the SWS 42 allows diagnostics to be executed without interrupting on-line data transfers.

The switch control module 86 (SCM) contains the components necessary to provide operational control of the SWS 42. The SCM 86 receives commands from and transmits status to the control subsystem 40 (CNS) through the MSL control LAN 95. The LAN interface unit 310 in the SCM 86 provides the SWS 42 interface to the MSL control LAN on line 95. This component is a commercially available off-the-shelf product. The switch controller 312 receives high level commands from the CNS 40 through interface 310 and converts these commands into detailed actions required by the SWS 42 components. Switch controller 312 is a commercially available controller product. The control interface unit 314 provides an interface between the commercially available switch controller unit 312 and the diagnostic distribution module 88. The control interface unit 314 is a commercially available product that contains the switch controller component interface. The control interface 314 does not contain enough I/O channels to support all of the command/status interfaces to each SWS 42 component. Therefore, the command/status information is passed from the control interface 314 via one parallel interface to the distribution unit 294 in DDM 88. The distribution unit 294 has sufficient I/O to support all command/status interfaces to each SWS 42 component.

The details of the drive subsystem 48 (DRS) are shown in FIG. 8. DRS 48 contains a recorder module 96 (RM) and a drive controller module 94 (DCM). The RM 96 maybe of various types but in the preferred embodiment is a rotary head recorder. The drive controller module 94 (DCM) provides the interface to the switch subsystem 42. This interface, used in conjunction with the internal logic of the DCM 94 and an interface which the DCM 94 provides to the RM 96, allows the transfer of file data and control information between the RM 96 and an interface subsystem 14, 16, 18 or 21. (See FIG. 2). The DCM 94 also provides an interface to the MSL control LAN from CNS 40 on line 95, which allows controlling the RM 96 which is not logically connected to an IFS 14, 16, 18 or 21.

The DCM 94 includes an interface component 316, a micro-controller unit 318 (MCC), a buffer unit 320, and a recorder interface 322.

The interface component 316 is a unit for interface with the SWS 42. It is designed for use in conjunction with the high speed parallel interface and allows the transfer of data between the IFS tape server, 14, 16, 18,

16

or 21, and the DCM 94 at burst rates of up to 640 Mbps. Interface unit 316 contains the interface assembly 324 which provides for the high speed parallel operation as well as a generic synchronous interface for the remainder of the DCM 94.

The MIC component 318 provides operational control for the components of the DCM 94, except for the interface unit 316, and the RM 96. The MIC 318 comprises a controller 326 and a LAN interface unit 328. The controller unit 326 is implemented with a 16 MHz microprocessor and associated control circuitry, including an interrupt handler, dedicated random access memory (RAM), and read only memory (ROM), address decoding, and bus arbitrator.

The controller unit 326 also provides a set of command registers through which the microprocessor interfaces with the interface unit 316, buffer unit 320, and recorder interface 322. It also includes a connection with the LAN interface unit 328.

The LAN interface unit 328 provides for communication between an external subsystem controller (CNS 40) and the controller 326 via the MSL control line 95. The LAN interface unit 328 is implemented with a LAN chip-set to provide the standard interconnection. The MIC unit 318 allows the DCM 94 to operate in either of two modes, on-line and off-line.

The DCM 94 enters the on-line mode when a logical connection is made between an IFS 14, 16, 18, or 21 and the DCM 94 through the switch subsystem 42. On-line operations occur when a particular IFS requires the use of a specific RM 96 to which the DCM 94 is connected for storing or retrieving file data. The MIC 318 receives commands from an IFS via the interface unit 316 in this mode and provides operational control accordingly. The received command requests an operational function such as data retrieval, data storage, file search, and the like. The operational control provided includes generating high level commands for each of the components and the RM 96, monitoring overall modular operations, and coordinating inter-component activities. It also includes collecting individual component and RM 96 status messages and command responses and assembling these messages and responses in a format suitable for transmission to the IFS requesting the function.

The DCM 94 enters the off-line mode when there is no logical connection between it and any IFS computer. In the off-line mode, the CNS 40, discussed in relation to FIG. 2, monitors the operability of both the DCM 94 and RM 96. It allows the IFS computer to maximize the use of the RM 96 to which the DCM 94 is connected. The MIC unit 318 receives commands from either the CNS 40 or IFS computer 14, 16, 18 or 21 via the LAN interface unit 328 in this mode and provides operational control based upon the command. The received command may instruct the performance of any operational capability, except the transmission of data to or from an IFS computer, since there is no logical connection. Off-line commands fall into one of two categories: diagnostic testing or preparation for on-line activity. The diagnostic testing commands allow the CNS 40 to monitor the operability of both the DCM 94 and RM 96. These commands are used to initiate the execution of the built-in test equipment available in both the DCM 94 and the RM 96. The commands preparing for an on-line activity allow an IFS computer to maximize the use of the RM 96. The commands enable activities, such as tape positioning before a logical connection is made between an IFS computer 14, 16, 18 or 21

5,412,791

17

and a RM 96, so that less post-connection time is lost due to searching for specific files on a tape.

The performance of the MIC 318 is obtained with a 16 MHz microprocessor implemented in the controller unit 326. The performance of the MIC 318 in the off-line mode is triggered by the LAN interface unit 328, which provides a serial communication path between the DCM 94 and the CNS 40.

The buffer component 320 compensates for both effective and burst data transfer rate disparities between the IFS tape server 14, 16, 18, or 21 and the RM 96. The buffer component 320 also provides verification of data being written to the tape and comprises a multi-access buffer unit 330, a buffer control unit 332, and a read-after-write compare unit 334.

The multi-access buffer unit 330 (MAB) provides the solid state memory necessary for temporary data storage as it is transferred between the individual IFS tape server and the RM 96. This memory off-loads some of the competition for the tape server's memory and allows portions of the file to be reaccessed in real-time to verify that it was actually written to tape and for re-writing it to tape, if necessary. The buffer control unit 332 provides addressing and control information to the MAB 330 and allows the MIC unit 318 to access the MAB 330 for use in diagnostic testing.

The read-after-write compare unit 334 (RWC) provides the logic necessary verify, during the write process, that the data is correctly, in a logical sense, being written to a tape. The RWC unit 334 also includes a means to indicate to the buffer control unit 332 that a portion of the file should be rewritten.

The MAB unit 330 is a 500 Mbit circular memory using commercially available 1 Mbit dynamic RAMs. The buffer size is that amount of memory which is necessary to compensate for the transfer rate variances and the expected latency from the time the data is written to tape until it is returned for verification. The MAB unit 330 provides a data interface to the interface unit 316, buffer control unit 332, the RWC unit 334, and the recorder interface unit 322. A combination of these interfaces allow simultaneous, non-conflicting access to the MAB 330, depending upon the type of command being executed. During a store to tape operation, the interface unit 316 transfers data to the MAB 330 while the recorder interface component unit 322 and the RWC unit 334 function simultaneously to receive data from different locations in the MAB 330. During a retrieve from tape operation, the recorder interface unit 322 transfers data to the MAB 330 while the interface unit 316 simultaneously retrieves data from a different location in the MAB 330.

During a diagnostic operation, the MIC unit 318 transfers data to the MAB 330 via the buffer control unit 332. The MAB 330 subsequently transfers the data simultaneously to the recorder interface unit 322 and, at an address lag, to the RWC unit 334 to verify operation of the DCM 94 and/or RM 96. Data is stored in the MAB 330 in a format which is compatible with the interface to the interface unit 316. Format conversions are performed by the other units to which the MAB 330 interfaces.

The buffer control unit 332 has a state feedback controller and the logic necessary to provide functional control of the memory in the MAB 330, such as addressing, interface arbitration, transfer direction indication, and the like. The buffer control unit 332 contains registers addressable by the MIC unit 318 which allow for

18

the transfer of high-level command and status messages. It also provides the MIC unit 318 with a data interface to the MAB unit 330 for use in diagnostic testing and provides an interface with the RWC unit 334 for determining when a portion of a file should be rewritten to tape. The buffer control unit 332 provides real-time control of the entire buffer unit 320 and its interaction with the MIC unit 318 is only on a high-level command/status transfer basis. The buffer control unit 332 uses well known prior art techniques which allow multiple interfaces to have simultaneous, contention-free access to a common memory. This technique is based on time-shifted memory data and address information for each of the interfaces supported.

The buffer control unit 332 also includes provisions for insuring that the file data is retrievable in a properly ordered manner. To do this, the buffer control unit 332 adds an identification marking to the data blocks before writing a block to tape. This is necessary because of the rewrite capability included in the implementation. When a file is read from tape, it is possible that duplicate blocks will be reproduced. This happens if an error occurs when the file was originally written, causing one or more blocks to be rewritten. The buffer control unit 332 uses the identification marking to determine which one of the duplicated blocks it should transfer to the interface unit 316 during a file retrieval operation. This capability insures the ability to retrieve error-free files in the order in which they were stored.

The RWC unit 334 uses a comparative circuit to provide positive verification that data was written to tape. This allows the comparison of data written to tape in diagnostic testing and provides a means by which defective tape is logically "cut-out" to insure that an acceptable bit-error-rate (BER) is maintained. The RWC unit 334 provides interfaces with both the MAB unit 330 and the recorder interface unit 322 to receive information previously written to tape. The RWC unit 334 performs a bit-for-bit compare operation to insure that the data received from the RM 96 is identical to that which was transferred to it for writing to tape. Detection of errors during this process indicates that the correction capability of the error-detection-and-code (EDAC) unit in the RM 96 was exceeded during the after-write reading of the data. This, in turn, implies an unacceptable imperfection in that tape section and causes the RWC unit 334 to instruct the buffer control unit 332 to rewrite the erroneous data on another section of the tape. By logically "cutting-out" bad tape sections in this manner, the uncorrected BER of the tape is effectively set without pre-screening it. This process, together with the tape's performance, and the error-detection-and-code unit of RM 96 implements an error correction strategy which insures an acceptable BER for data stored to and retrieved from tape. The buffer unit 320 fully compensates, without degrading the effective channel transfer rate, for the difference between the fixed data transfer rate of RM 96 of 240 Mbps and an effective channel transfer rate which varies between 1 Mbps and the RM 96 (240 Mbps) rate.

The recorder interface unit 322 (RCI) compensates for the data format and the logic differences between the MAB 330 and the RM 96. The RCI 322 also provides a means for converting the format and logic of the commands and status responses transferred between the MIC unit 318 and the RM 96. The RCI 322 is implemented using a data interface unit 336 and a control interface unit 338.

19

5,412,791

20

The data interface unit 336 (DIF) contains the balanced ECL (emitter coupled logic) line drivers and receivers required to transfer data between the DCM 94 and RM 96. The DIF 336 provides a unique interface pair (one input, one output) to both the MAB unit 330 and the RM 96 to enable this transfer. The design of the DIF 336 insures the ability of the RCI 322 components to simultaneously transfer data to and from the RM 96 thereby supporting the execution of read-after-write verification by the DCM 94. The DIF 336 also provides for clocking differences between the DCM 94 and RM 96 and provides an interface with the control interface unit 338 for receiving mode information. The DIF 336 operates in one of two modes based upon this information, normal and loop-back. In the normal mode, the DIF 336 allows for data transfers between the MAB unit 330 and the RM 96. In the loop-back mode, which is entered only during diagnostic testing, the DIF 336 accepts data from the MAB 330 and transmits this data directly to the RWC unit 334 for comparison with the original data to verify the DIF 336 operation.

The control interface unit 338 (CIF) contains a set of command registers through which the MIC unit 318 components communicates with the RM 96. The CIF 338 also includes the ECL interfaces over which the auxiliary data and the error flag are transmitted and the interface over which RM 96 commands and the status responses are transmitted.

The error flag on line 340 is a signal that indicates some level, less than maximum, of error correction capability was exceeded by the error-detection-and-correction (EDAC) of the RM 96 during a file retrieval. This signal determines when, on initial data write, to re-write a data block, or on subsequent data reads, when a tape has degraded to a point that warrants re-writing it to another tape. This signal is used to monitor BER degradation over repeated tape use and time. The level at which the signal is asserted is programmable. The signal is also monitored the first time a file is written to tape to insure that it is written within a certain margin of correctability. Data blocks which fail the bit-for-bit compare process will be rewritten on the tape. It is possible, that a file is written to tape such that it passes the bit-for-bit compare during the read-after-write but the error correction capability is stressed to its limit in providing the correct data as output. In such case, a lower level will be programmed for which the error flag will signal that some correction capability has been exceeded.

The auxiliary data interface lines 342 provide a means to store housekeeping information on the cassette tape's longitudinal tracks for later use in the file search or retrieval processes. The auxiliary data interface lines 342 are implemented as a unique pair (one input, one output) of serial interfaces. The control interface 344 provides a bidirectional, serial interface which RM 96 commands and through which the status messages may be transferred.

The data interfaces provided by the DIF unit 336 transfer data independently, and simultaneously, if desired, at a fixed rate of the RM 96 (240 Mbps) to and/or from the RM 96. The CIF unit 338 transfers information and status to and from the RM 96 at a standard rate of 9.6 Kbps. The auxiliary data interface 342 transfers information between the RM 96 and the CIF unit 338 at a rate of 38.4 Kbps.

The recorder module 96 provides a rotary head recorder which is known in the prior art and which is

used for permanent storage and subsequent retrieval of file data in the mass storage system. The rotary recorder employs four pair of heads (one record, one reproduce), and uses a helical scan technique to provide high speed file storage and retrieval. The recorder uses the cassette tape as a storage medium. It has the capability of transferring data rates at 200 Mbps or higher, and is also capable of obtaining high tape speeds (non-record/reproduce) to support file searches. The RM 96 implements an innovative error correction strategy to insure an acceptable BER. Further, it includes a sophisticated set of built-in test features to insure unit testability. It also employs slant azimuth recording which provides a significantly better tolerance to tracking errors than non-slant azimuth recording and satisfies the BER requirement. Since the magnetic signals on adjacent tracks are perpendicularly oriented, much less induced noise is experienced due to magnetic coupling between a reproduced head and the magnetic signals from tracks adjacent to the track actually being read when tracking errors occur. The use of slant azimuth recording is an effective means to compensate for tracking errors due to the environmental effects on the tapes. It is also an effective means to ensure crossplay, since it compensates for the tracking errors due to mechanical tolerances in the construction of different recorders.

The RM 96 provides a parallel interface capable of transferring data at a fixed burst rate of 240 Mbps. The interface is provided as a pair (one input, one output) of uni-directional interfaces each of which are capable of obtaining the 240 Mbps transfer rate. Each interface is implemented with electronics to provide independent data transfers at 240 Mbps to support the read-after-write verification function. The RM 96 can transfer or receive contiguous blocks of information at the burst transfer rate based upon the issuance of a single command.

The recorder's tape transfer drive motors support the capability to perform fast file searches by providing the ability to shuttle the tape at speeds of up to 300 inches per second (IPS) in either the forward or reverse directions. The motors also provide smooth acceleration in transitioning between tape speeds to prevent tape damage. An auto load function includes automatically sensing a tape cassette size, loading the cassette on the reel hubs, and wrapping the tape around the helical scanner. The recorder 96 performs the complete autoloading function within a maximum of three seconds after the cassette has been inserted approximately five inches into the recorder by either a robotic accessor or manually by a human operator. The recorder 96 performs an unload function with similar but reversed steps upon command from the DCM 94. The execution of this command results in a partial ejection of the cassette (approximately three inches) which is sufficient for removal by a robotic accessor. As indicated earlier, the recorder module 96, as described, is available commercially.

The storage subsystem 104 (FIG. 2B) (STS) implementation of Global Cassette Access includes rotational modules 112 (RTM). The RTM 112 is implemented as a fixed rack which has dimensions of 40" x 40" x 40". The RTM 112 is configured with individual storage slots for the cassettes such that it uniquely stores one of two cassette sizes. The first configuration of the RTM 112 is illustrated in FIG. 9. In this configuration, the RTM 112 provides two columns, 346 and 348, of twenty storage slots in a side-by-side manner on each of two sides. A side is defined as a module face which normally makes

5,412,791

21

cassettes accessible to a CAM 106 (see FIG. 2), as opposed to an end which is a face that does not normally make cassettes accessible. Each storage slot 350 houses a single cassette, such that this configuration contains a total of 80 cassettes (40 cassettes accessible on each side).

In the second configuration shown in FIG. 10, the RTM 112 provides a single column 352 of 20 storage slots 354 on each of its two sides. Each slot 354 in this configuration is capable of accepting a single larger size cassette such that it contains a total of 40 cassettes (20 accessible on each side). The storage slots 350 or 354 are mounted inside the RTM 112 on a base 356 which is capable of rotation. This base 356 may be rotated so that all of the cassettes accessible on one side of RTM 112 are made accessible on the opposite side. The RTM 112 receives rotation commands from the CNS 40 through a well known type of interface that it provides to the storage control LAN signal on line 110 (FIG. 2B) from CNS 40. The storage control LAN signal on line 110 uses an Ethernet type LAN to which all the RTM 112s and all the transport subsystem 56 modules (FIGS. 1 and 2B) are connected for receiving commands from and transmitting status to the CNS 40.

The RTM 112 uses precision stepper motors to provide the rotation capability. Rotation is capable in 90 degree increments and is provided in a smooth manner such that the cassettes are not shaken excessively. Each storage slot 350 and 354 implements a well known retention mechanism for retaining the cassettes.

As can be seen in FIG. 11, which is a proposed floor plan of the Mass Storage Library of the present invention, the RTM 112 is designed to attach to a track component (not shown) along with other RTM modules 112 to form a storage row. The RTM 112 design allows for module stacking. This enables the configuration of storage columns which comprise up to four vertically centered RTM 112s.

In the preferred embodiment of the rotational module 112, some media storage slots 350 are used for housing the recorder module 96, as shown in FIG. 9. The mounting provisions are such that the recording module 96 provides a cassette interface with the cassette accessor module 106. The ability of the RTM 112 to rotate provides the system with a built-in redundancy feature. The RTM 112 is rotated in the event of either a cassette accessor module 106 failure or a recording module 96 failure to allow a cassette to be accessed by a different cassette accessor module 106. This redundancy is a key element in the design.

Each of the bulk load modules 358 (BLM) (see FIG. 2) serves as a cart for carrying cassettes which are to be loaded or have been unloaded from the lateral transfer modules 108. The BLM 358 is implemented as a rollable frame with provision for mounting two cassettes storage units, each of which is a box like container with individual storage slots so that it stores cassette tapes. In one configuration, the cassette storage unit provides two columns of ten storage slots in a side-by-side manner. The total configuration contains twenty cassettes. The storage units store the cassettes such that they are accessible from only one side of the cassette storage unit. The bulk load module 358 is easily moved by one person over tile or carpeted floors and up and down ramps. A fully loaded cassette storage unit weighs approximately 30 pounds and is easily removed from the BLM 358 and inserted into the lateral transfer module 108 by one person. This allows a quick means of manu-

22

ally loading a large number of cassettes in the lateral transfer module 108 without reducing its full operational state.

The transport subsystem 56 (see FIG. 2) is designed to insure that each cassette can be made accessible to each cassette accessor module 106 and can be used in conjunction with every recorder module 96. Each cassette accessor module 106 (CAM) has an accessor control electronics component which serves as the control function for the individual manipulator assembly. The control electronics receives the manipulator commands from the control subsystem 40 on line 110 in FIG. 2. These commands are converted into discrete manipulator movement instructions which are transmitted to the manipulator assembly through a modulated duplex infrared communication link. Likewise, the electronics receives manipulator status through the infrared communication link and performs the conversion necessary for transmitting the status to the control subsystem 40 through the communications line 110. The control electronics 357 are mounted in the rotational module 112. See FIGS. 9 and 10. By placing this module in the middle of a storage row's length, the manipulator assembly's communication range can be doubled. Use of the control electronics module 357 to provide manipulator assembly communications eliminates the possibility of physical cable interference that might otherwise be associated with movement.

It will be noted that each cassette accessor module 106 moves between and parallel to two rows of rotational storage modules 112. Thus, the faces or sides of the rotational modules 112 which face inwardly towards the cassette accessor module 106 can be accessed by the cassette module 106. Thus, the CAM 106, which is located between rows 360 and 362 of rotational modules 112, can access anyone of those modules which are facing inwardly to the space between the rows 360 and 362. Assume, however, that the CAM 106 between rows 360 and 362 malfunctions. In that event, those rotational units 112 in row 362 can be rotated 180 degrees to enable the CAM 106 between rows 362 and 364 to service them. In like manner, the rotational modules 112 in row 360 can rotate 180 degrees to allow the CAM 106 between rows 360 and 366 to serve them. Thus, a natural redundancy is built into the system because of the rotational ability of the RTM 112.

It can also be seen clearly in FIG. 11 that the lateral transfer modules 108 are attached at the ends of the rotational modules 112, thus allowing any cassette in any one rotational module 112, in any given row, to be removed by the CAM 106 and transferred to a lateral transfer unit 108. Lateral transfer unit 108 will move that cassette from one lateral transfer unit to the other until it gets to the desired row of rotational modules 112 where the associated CAM 106 will remove it from the lateral transfer unit 108 and place it in the proper location in rotational unit 112.

Thus, there has been disclosed a novel mass storage library system which has three major functional subsystems, including a file server subsystem, a control function subsystem, and a media handling subsystem. The file server subsystem includes first a disk server interface computer which provides high-performance on-line archive capability. A plurality of second tape server interface computers couple the system users to any storage elements in the system through a switch module utilizing crossbar switches coupled to each other.

23

The media handling subsection provides global access to storage elements and includes storage element modules which allows accessor modules to access storage elements stored on either side of the storage module, thus providing built-in redundancy to the system. In addition, lateral transfer modules enable storage elements to be transferred from one row of rotational modules to another row.

The control subsection provides the allocation and de-allocation of all of the resources in addition to monitoring system activities, providing system diagnostics and providing system information in hard copy or soft copy format.

While the invention has been described in connection with a preferred embodiment, it is not intended to limit the scope of the invention to the particular form set forth, but, on the contrary, it is intended to cover such alternatives, modifications and equivalents as may be included within the spirit and scope of the invention as defined in the appended claims.

We claim:

1. A mass data storage and retrieval system comprising:

- a plurality of means for information storage forming a mass storage library;
- a data directory for maintaining a directory of the information stored in the mass storage library and for generating, in response to a request, a data location output signal identifying the location within the mass storage library of information responsive to the request;
- a plurality of data recorder modules for receiving and reading information from and writing information to a selected means for information storage;
- interface means bi-directionally coupled to the data recorder modules for simultaneous reading and writing of information from and to the selected means for information storage received by the data recorder module;
- a control computer coupled to the data directory for receiving the data location output signal and, in response thereto, generating a first command signal output to the mass storage library for selecting and loading the selected means for information storage in the data recorder module and for generating a second command signal for coupling the interface means to the recorder module loaded with the selected means for information storage.

2. The system as in claim 1 wherein the mass storage library comprises:

- a plurality of storage modules each having a plurality of openings in at least one side thereof for retaining means for information storage; and
- an accessor associated with the storage modules for retrieving one of the means for information storage from an opening in the storage module for use by one of the data recorder modules during simultaneous information read and write.

3. The system as in claim 2 further comprising a lateral transfer means for moving the means for information storage between adjacent storage modules.

4. The system as in claim 3 wherein the lateral transfer means comprises a storage unit movable on a trolley from one of the adjacent storage modules to another, said trolley having slots for retaining the means for information storage, the slots being accessible to each accessor to enable lateral transfer of the means for information storage between adjacent storage modules.

5,412,791

24

5. A data retrieval and recording system having a computer for interfacing with a recorder and a storage library to read data from or write data to a predetermined area of a selected storage medium, comprising:

- means for interfacing the computer to the recorder for simultaneous read and write operations; and
- controller means coupled to said means for interfacing for receiving operational function commands from the computer and for configuring the means for interfacing for simultaneous read and write operations, the controller means further coupled to the storage library for loading a selected storage medium in the recorder, and further coupled to the recorder for queuing the loaded storage medium to the predetermined area.

6. The system as in claim 5 wherein the means for interfacing comprises:

- a buffer memory; and
- a buffer control circuit coupled to the buffer memory for storing data written from the computer in one area of the buffer memory and for storing data read from the storage medium to be stored in another area of the buffer memory.

7. The system as in claim 6 further including a read-after-write means including means for reading and sorting data after writing to the storage medium in one area in the buffer, and means for storing the data to be recorded in another area, further including means for comparing the data to be recorded with the data read from the storage medium, and means responsive to the means for comparing for generating an error signal indicating that the data from the computer was recorded incorrectly.

8. The system as in claim 7 further including means responsive to the error signal from said means for generating for causing incorrectly recorded data to be read from the buffer and rewritten on the storage medium.

9. A mass data storage system, comprising:

- a storage module having means for retaining a plurality of recording media on a plurality of sides thereof;
- at least one recorder module having means for writing data to and reading data from a recording media;
- a recording medium accessor module for retrieving a retained recording medium from a side of the storage module and for inserting the recording medium in the recorder module; and

means for providing access to recording media retained on any side of the storage module for insertion in the recorder module.

10. The system as in claim 9 wherein the means for providing access comprises means for rotating the storage module to associate the recording media on the side of the storage module with the accessor module for retrieving the recording media by the accessor module for insertion into the recorder module.

11. The system as in claim 9 including a plurality of separate storage modules and accessor modules further comprising:

- at least one transfer module adjacent the plurality of storage modules having openings for retaining recording media therein deposited by an accessor module; and

means for moving said transfer module between individual ones of the plurality of storage modules to enable transfer of recording media from one stor-

5,412,791

25

age module to and from the transfer module by an accessor module for movement between modules.

12. A mass data storage library system, comprising: a plurality of data storage devices in a library each having areas for storing data;

means for identifying a designated storage device in the library for access and for identifying and selecting an area of the designated storage device on which data is to be read or written;

means for accessing the designated data storage device to read or write data in the selected area; and file directory means for storing file directory data representing the identity of the designated storage device, the area of that device selected for the data and a code identifying the data stored on the accessed storage device at the selected area.

13. The system as in claim 12 further comprising: means responsive to a request for selected data for accessing the file directory means to identify the storage device and the area thereon containing the selected data;

means for retrieving the data storage device having the selected data; and

a record/playback module for selectively receiving the data storage device having the selected data.

14. A system for providing any one of a plurality of host computers with access to stored data, comprising: a mass storage data library comprising a plurality of individually accessible data storage devices;

a data directory archive for maintaining an address directory of the data stored on each data storage device;

a plurality of data record/playback modules, each of said modules reading data from and writing data to a selected data storage device;

interface means for bi-directionally coupling a host computer to one of the data record/playback modules; and

a controller computer coupled to the interface means, the data directory archive and the mass storage library for receiving data access requests and a directory address for the requested data and, in response thereto, generating a first signal causing a data storage device containing the requested data to be loaded into a selected data record/playback module and a second signal for configuring the interface means to couple the requested data to a requesting host computer.

15. A system for providing access by means of a user interface to a mass data storage and retrieval system, comprising:

means for storing data in selected locations in a storage medium;

means for maintaining a file directory archive identifying the location of all stored data in said storage medium;

means for coupling a computer to the user interface for receiving requests to access data on said storage medium; and

means for coupling the computer to the file directory archive and said storage medium and for utilizing the data location identified by said file directory archive to retrieve requested data.

16. A system for mass storage of data, comprising: means for storing recording media in slots on at least two sides of a storage module;

26

means for coupling an access module to said storage module to retrieve and return recording media from and to said slots; and

means for enabling said access module to retrieve and return recording media to slots on either of said at least two sides of said storage module.

17. The system as in claim 16 wherein the means for enabling comprises means for rotating the storage module to expose each side of the storage module to the access module.

18. A mass data storage system including: a plurality of data storage modules;

a data directory archive for maintaining a directory of the data contained on each data storage module;

data record/playback modules for receiving a selected data storage module;

a plurality of interface computers coupled to a plurality of host computers;

a file directory coupled to the interface computers and the data directory archive for receiving a data request to locate the data storage module containing the requested data in the data directory archive, and for generating a data address location output signal;

a control computer for receiving the data address location output signal and, in response thereto, generating a signal for loading the data storage module containing the requested data into a selected data record/playback module; and

an input/output data channel for coupling the interface computer to the selected data record/playback module to access the stored data on the data storage module.

19. The system in claim 18 wherein the input/output data channel further includes means responsive to an access request signal for selectively coupling any of the plurality of interface computers to any of the data record/playback modules.

20. The system is in claim 18 wherein the input/output data channel includes means for simultaneously carrying data to and from the selected data record/playback module.

21. A system for interconnecting a plurality of data processing means, comprising:

a switch module having a plurality of input/output ports, each port connected to one of the plurality of data processing means; and

a switch control module responsive to a command signal for generating a control signal output to the switch module for coupling any two selected input/output ports together to form a signal path facilitating data communications between any two of the plurality of data processing means.

22. The system as in claim 21 further including a diagnostic module coupled to the switch control module for generating a diagnostic signal output to the switch module to execute a diagnostic routine for testing the switch module.

23. A mass data storage and retrieval system comprising:

a plurality of means for information storage forming a mass storage library;

a data directory for maintaining a directory of the information stored in the mass storage library and for generating, in response to a request, a data location output signal identifying the location within the mass storage library of information responsive to the request;

5,412,791

27

a plurality of data recorder modules for receiving and reading information from and writing information to a selected means for information storage;
 interface means bi-directionally coupled to the data recorder modules for reading and writing of information from and to the selected means for information storage received by the data recorder module;
 buffer means for storing information transmitted to the data recorder module as written to the means for information storage and storing information received from the data recorder module as read from the means for information storage; and
 a control computer coupled to the data directory for receiving the data location output signal and, in response thereto, generating a first command signal output to the mass storage library for selecting and loading the selected means for information storage in a data recorder module and for generating a second command signal for coupling the interface means to the recorder module loaded with the selected means for information storage.

24. The system as in claim 23 wherein the buffer means further includes read-after-write comparison means for comparing the stored read information with the stored written information to verify the accuracy of the recording.

28

25. A system for interconnecting a plurality of data processing means, comprising:
 a switch module having a plurality of input/output ports, each port connected to one of the plurality of data processing means;
 means, responsive to a request for interconnecting two data processing means, for generating a command signal identifying the data processing means to be interconnected; and
 a switch control module responsive to the command signal for generating a control signal output to the switch module for coupling two selected input/output ports together to form a signal path for data communications between the identified data processing means.

26. The system as in claim 25 wherein the switch module comprises a plurality of switch means in a matrix configuration serially connected in response to the control signal to form the signal path coupling the two selected input/output ports.

27. The system as in claim 26 wherein the plurality of switch means comprise a plurality of crossbar switches.

28. The system as in claim 25 wherein the data processing means include a computer and data record/playback devices, the switch module interconnecting the computer and data record/playback devices to provide for computer access to the data record/playback devices.

* * * * *